Entity Suggestion by Example using a Conceptual Taxonomy

Yi Zhang*, Yanghua Xiao†, Seung-won Hwang‡, Haixun Wang§, X. Sean Wang¶, and Wei Wang‖
* † ¶ ‖ Fudan University, Shanghai, China,
* z_yi11@fudan.edu.cn
† shawyh@fudan.edu.cn
¶ xywangcs@fudan.edu.cn
‖ weiwang1@fudan.edu.cn
‡ Yonsei University, Seoul, South Korea
‡ seungwonh@yonsei.ac.kr
§ Facebook, Menlo Park, CA, USA
§ haixun@gmail.com

arXiv:1511.08996v1 [cs.IR] 29 Nov 2015

***Abstract*— Entity suggestion by example (ESbE) refers to a type of entity acquisition query in which a user provides a set of example entities as the query and obtains in return some entities that best complete the concept underlying the given query. Such entity acquisition queries can be useful in many applications such as related-entity recommendation and query expansion. A number of ESbE query processing solutions exist in the literature. However, they mostly build only on the idea of entity co-occurrences either in text or web lists, without taking advantage of the existence of many web-scale conceptual taxonomies that consist of hierarchical isA relationships between entity-concept pairs. This paper provides a query processing method based on the relevance models between entity sets and concepts. These relevance models can be used to obtain the fine-grained concepts implied by the query entity set, and the entities that belong to a given concept, thereby providing the entity suggestions. Extensive evaluations with real data sets show that the accuracy of the queries processed with this new method is significantly higher than that of existing solutions.**

I. Introduction

Entity suggestion by example (ESbE) has been widely investigated in different scenarios. In a typical scenario, a system accepts a set of example entities provided by a user as a query q, and retrieves a set of entities such that these entities, along with q, complete some concepts. For example, a user may type $\{$*China, India, Brazil*$\}$ as a query. It is quite possible that the user bears the concept *BRIC* in mind but cannot recall all of its members, so he/she enters these example entities of the concept for the purpose of retrieving the remaining ones. Then, the remaining entity *Russia* should be returned as the result. We give more examples in Table I. Entity suggestion by example is also known as *entity list completion* [1], [2], *entity retrieval* [3], [4], *entity recommendation* [5] or *entity query by example* [6], [7] in different settings.

In general, the example entities can also imply some other concepts such as *country* in our previous example. However, a general concept usually leads to many less related entities. For example the concept *(*country) might mislead to the less related entities such as *US*. Hence, in ESbE, we aim to find fine-grained (specific) underlying concepts so that we can suggest most semantically related entities. We will elaborate the selection of concepts in Section V.

ESbE has many applications. For example, in a search engine, a good recommendation of *'related entities'* [8] can be useful in encouraging users to click on the links to these entities [9]. As another example, in a spreadsheet application, a tool that can automatically populate a spreadsheet or list can save on manual labor. For instance, when a user wants to create a table of BRIC countries in the spread-sheet, it may reduce his/her effort significantly if the table can be automatically populated after a few examples are entered. In general, ESbE can be viewed as a type of query by example [10], and can be useful in some question answering systems [11]. With more and more knowledge bases published, entity information becomes increasingly abundant and queries by exemplar entities on knowledge bases will become more popular [12], [13].

Weakness of Previous Approaches: Many solutions to ESbE have been developed. These solutions can be classified into the following three categories. The solutions in the first category tend to use co-occurrence as the basic recommendation mechanism. A well-known example is Google Set. The basic idea is to recommend the entities that most frequently co-occur with the example entities. The solutions in the second category assume that the query set belongs to some lists and estimate how likely it is that each item belongs to a list containing the query items. An example in this category is SEISA [14]. The solutions in the third category rank all the entities based on how much their properties overlap with those of the example entities, and take the highest ranked entities not already in the query as the result.

However, these solutions, being unaware of the concepts, especially fine-grained concepts underlying the example entities, suffer from the following limitations:

1) First, *co-occurrence does not necessarily imply conceptual coherence* [15] *between the recommended entities and the example ones*. Best result entities are often those that share the same concepts with the example entities. For example, when a user searches with $\{$*China, India, Brazil*$\}$, the intention is very

TABLE I: Entities and Possible Fine-grained Underlying Concepts

Entity List	A Suggested Entity	A Possible Fine-grained Underlying Concept
China, India, Brazil	Russia	BRIC
Alibaba, Tecent	Baidu	BAT
swimming, marathon	bicycle ride	Ironman Triathlon
Islam, Buddhism	Christianity	The three major religions
Standard Poor's, Moody's	Fitch Group	Big Three(credit rating agency)
Roger Federer, Rafael Nadal, Novak Djokovic	Andy Murray	Big Four(tennis)

likely to find the *BRIC countries*. In general, *Russia* is a good entity to be recommended since entity *Russia* is better in keeping the concept coherent than entity *USA* is even if it frequently co-occurs with all the example entities respectively. Note that it is unrealistic to count the number of co-occurrence among a set of entities in advance by enumerating all possible combinations of entities even in an off-line procedure.

2) Second, *the membership of entities in lists may be too weak a signal to be used to recommend the right entity.* In the state-of-the-art methods for set expansion SEAL [16] and SEISA [14], web lists are used to construct an entity-list-membership network and candidate entities are ranked by a random walk based measure. In general, one should recommend an entity more strongly if it can be reached with higher probability by random walk starting from the given entities. However, because of the weak membership between entities and lists, the random walk along the membership edges is likely to find noisy entities. Even when there are lists containing all query examples, we cannot decide which is more desirable without knowing the concept behind it. They may refer to a concept too general to introduce false positives, or even they are lists randomly involving some entities from the web. For example, given {*China, India, Brazil*}, a very typical list they belong to may be a list of many countries. Then some other common entities in this kind of lists such as *USA, Britain* etc. may be ranked high in the results. These typical developed countries should obviously not be suggested first when given some of the BRIC countries.
3) Third, *considering only the overlap of the properties of the example entities and those of the candidate entities does not always lead to the correct set expansion.* For example, given three entities {*Industrial and Commercial Bank of China, China Construction Bank, Bank of China*}, though the entity *China Merchants Bank* shares with most of the query entities in terms of its properties such as *location, owner* and so on, this entity may not be a good answer. The reason is because *Big Four State-owned Bank of China* may be the most relevant and fine-grained concept behind the three given banks while this concept does not cover *China Merchants Bank*. (The other Big-Four is *Agriculture Bank of China*.)

Of course, entity suggestion can be quite subjective, and the best entities to recommend may vary from person to person. Our work is to recommend entities by determining the fine-grained concepts underlying the query examples. Thus, we will give priority to the entities related to a fine-grained underlying concept which we think is a general sense for a generic user. With enough contextual information, our method may be extended to provide personalized entity suggestion, but this extension is beyond the scope of this paper.

Advantages of Conceptual Taxonomies: Recently, many web-scale conceptual taxonomies consisting of isA relationships between entity-concept pairs, such as Microsoft's Probase and Google's isA database, have become available. These knowledge bases are constructed by Hearst patterns from a web corpus. The abundant concept information in these knowledge bases brings us new opportunities to process ESbE queries.

1) First, *the taxonomy allows us to explicitly and accurately model the concepts of example entities as well as their complicated relationships.* As we have seen in our examples above, the concept modeling is critical for ESbE. A large scale conceptual taxonomy allows us to explicitly represent the semantic of an entity or an entity list by its concept distributions.
2) Second, *the frequency information in taxonomy makes the inference more accurate*. Some taxonomies, such as Probase, have the frequency information for the isA pairs. This allows the computation of many desired metrics, such as the typicality of an entity/concept. For example, *emerging market* is more typical than *country* as a concept for the example entities {*China, India, Brazil*}. Using this more typical concept allows us to infer *Russia* instead of some other country entities.
3) Third, *the hierarchical structure of the conceptual taxonomies enables more accurate inference of the result entity.* Note that the hierarchy can be used to estimate the specificness of a concept. For example, *China* is*A country*, and also a *developing country*. Based on the taxonomies, we also know that a *developing country* is also a *country*. Thus, *developing country* is a relatively more specific concept. Using this concept allows to suggest more related entities by reducing the possibilities of false positives.

Challenges: In this paper, we propose two probabilistic approaches to infer the entities that belong to the concepts implied by the exemplar entities. Our approaches sufficiently utilize the concept information of entities for inference. To do so, we face the following challenges:

First, *how to aggregate the conceptual information of*

example entities without introducing extra noise is still challenging. To see this, we show that some naive aggregations, such as finding the *least common ancestor* (LCA) concept of query entities, fail in our running example. The LCA concept of {*China, India, Brazil*} in Probase consists of concepts such as *country, nation* and so on. Thus, it is very likely to return some entities such as *USA* by finding other entities of the LCA concepts such as *country*. To overcome this challenge, we propose a Bayesian approach and a Noisy-Or model to evaluate the relatedness between the concepts and the examples to represent the semantic distribution for the example entities in our solution.

Second, *how to determine fine-grained concepts instead of concepts too general from the concepts related to the query examples is difficult.* General concepts may introduce more false positive entities than fine-grained concepts may, thus we propose two kinds of methods to solve this problem. The first one is directly penalizing popular concepts which is an entity based approach, and the second one is a granularity-aware approach which is based on the hierarchical structure of the taxonomies.

Third, even if desired fine-grained concepts can be found, a false positive may still be generated without careful treatment. Since isA relationships between entities and concepts are extracted from a large corpus by patterns, some false positive entities may already exist in the list of the entities of a concept. For example, *China* is mentioned as a developed country in some corpus, thus it appears in the entity list of *developed country* as well. However, by a careful examination, one may notice that *China* is more closely related to the concept *developing country* in the taxonomy. To address this challenge, we propose a probabilistic relevance model (in Section IV-A), which leverages the typicality of entities to every related concept to make the inference bias towards the promising entities. We also model our problem as an optimization problem, which minimizes the difference before and after the admission of a candidate entity (in Section IV-B). The optimization model leads us to promising entities.

The rest of the paper is organized as follows. In section II, we review the related work. In section III, we introduce some background knowledge and present the baseline method. In section IV, we propose two probabilistic models for our problem. Section V elaborates how we compute the models. We present the experimental study in Section VI and conclude the paper with Section VII.

II. Related Work

In this section, we review the related work, which can be classified into three categories: *related entity recommendation*, *entity set expansion* and *short text conceptualization*.

Entity Recommendation

Related entity recommendation can be categorized into the following two categories: First, to *recommend related entities for search assistance*, Blanco et al. [17] proposed a recommendation engine Spark to link a user's query word to an entity within a knowledge base and recommend a ranked list of the related entities. To guide user exploration of recommended entities, they also proposed a series of features to characterize the relatedness between the query entity and the related entities. Unlike this work assuming a single entity as a query, our work recommends related entities for a set of query entities of the same hidden concepts. Steffen et al. [18] proposed a similar entity search considering diversity, however, in our work, we assume an entity linked to related and fine-grained concepts more important.

Second, for *query assistance for knowledge graphs*, GQBE [19] and exemplar queries [12] studied how to retrieve entities from a knowledge base by specifying example entities. For example, the input entity pair {Jerry Yang, Yahoo!} would help retrieve answer pairs such as {Sergey Brin, Google}. Both of them projected the example entities onto the RDF knowledge graph to discover result entities as well as the relationships around them. GQBE used the weighted graph as the underlying model and exemplar queries used the edge-weighted one. All these works used the subgraph isomorphism as the basic matching scheme, which in general is costly. Our objective is to infer entities that preserve the semantic of the example entities. All the example entities generally share the same type.

Entity Set Expansion

The goal of this line is, given a set of seed entities, to discover other entities in the same concept. Google Sets [20] is a product implementation used to populate a spreadsheet after users provide some instances as examples. Inspired by Google Sets, many research work followed [21], [14], [22], [23], [24], to measure the membership strength of an item for a hidden concept exemplified by query entities.

However, this work assuming a single concept (always a too general concept) to explain the query set, cannot disambiguate when the query set conceptualizes to multiple fine-grained concepts, such as a camera brand and Japanese company. Our work assumes a query set can bear multiple fine-grained concepts, and aggregates a concept distribution to accurately infer related entities that reflect all the related concepts.

Related problems include semantic search tasks studied in Bron et al. [7], of taking example instances and the textual representation of a relation, to complete the list of examples. Another example is harvesting tables on the Web, and retrieving the table that completes the example instances and description [14], [25]. Compared to this work, our task is more challenging relying solely on examples without explicit description of a relation or table.

Text Conceptualization

Conceptualization aims to map a short text to a set of concepts as a mechanism of understanding text. Lee et al. [26] proposed context-dependent conceptualization to capture the semantic relations between words by combining Latent Dirichlet Allocation with Probase. Song et al. [27] developed a Bayesian inference mechanism to conceptualize words and short texts. The ultimate objective of conceptualization is to find the concepts that best capture the semantic of the short texts. However, conceptualization combines all related concepts together without considering the semantic granularity of the concepts which increases the risk of recommending

false-positives. Our work is to find fine-grained concepts underlying the query, and the next step of entity inference, of identifying the most related entity, is beyond the scope of conceptualization as well.

III. Background and Baseline Approach

In this section, we briefly review the conceptual taxonomy Probase upon which our solutions are built. We also argue that the straightforward solution based on Probase and the baseline solution based on the k-nearest neighbor algorithm cannot solve our problem.

A. Probase and isA relationships

We use a web-scale conceptual taxonomy for the inference of the suggested entities. Probase [28] is a universal, general-purpose, probabilistic taxonomy automatically constructed from a corpus of 1.6 billion web pages. Probase contains 2.7 million concepts and 4.6 million isA (a.k.a., hypernym-hyponym) relationships among the concepts/entities, which is suitable for us to describe the query entities. Each isA relationship saying (e isA c) is associated with the frequency ($n(e, c)$) that e isA c is observed from the corpus. The frequency allows us to compute the *typicality* of e under concept c, i.e., $P(e|c)$, which can also be interpreted as the probability that e is an instance of c.

Given Probase, a baseline of set expansion is to find the *least common ancestor* (LCA) concept of the query entities, then return an instance with the highest typicality for the LCA concept identified. However, it does not perform well, as the LCA concept may not exist in Probase. Furthermore, even the desired concept is identified, entity inference remains as a challenge, as typicality does not reflect relationships with query entities.

B. A KNN based Approach

As another baseline, we consider a k-nearest neighbor (KNN) based solution of returning the top-k entities with the highest $sim(q, e)$.

Thus, the key is defining sim using the concept distribution of entities derived from Probase. Specifically, given the query entity set $q = \{e_1, e_2, ..., e_n\}$, each entity e is associated with a concept vector $\mathbf{C}(e) = \{< c_i, P(e|c_i) >\}$, where $P(e|c_i)$ is the typicality of the instance e given the concept c_i. Given Probase, $P(e|c_i)$ can be computed by the following equation:

$$P(e|c_i) = \frac{n(e, c_i)}{n(c_i)} \tag{1}$$

where $n(c_i)$ is the number of occurrences of the concept c_i in Probase, and $n(e, c_i)$ is the number of occurrences of e as an instance of c_i.

Similarly, we can define the concept vector $\mathbf{C}(q) = \{< c_i, P(q|c_i) >\}$ for the query entity set q, such that each c_i is a concept of at least one entity in q and $P(q|c_i)$ is the typicality to observe any one entity in q under the concept c_i in Probase. $P(q|c_i)$ can be computed by the following equation:

$$P(q|c_i) = \frac{\sum_{e \in q} n(e, c_i)}{n(c_i)}$$

TABLE II: Notation Table

Notation	Description
E	the universal set of entities
q	the set of query entities
C	the universal set of concepts
$rel(q, e)$	the relevance of an entity e to q
$P(e\|c_i)$	the typicality of the entity e given the concept c_i
$P(c_i\|e)$	the typicality of the concept c_i given the entity e
$P(c_i\|q)$	the typicality of the concept c_i given a query set q
$P(c_i)$	the typicality of concept c_i
$\delta(c_i)$	the indicator of c_i being a fine-grained concept
$c(e)$	the concept set of the entity e
C_q^k	the set of k fine-grained concepts underlying the query q
$h(q_i\|c)$	the expected number of steps starting from q_i to c

Once concept vectors for q and e are defined, similarity can be computed as the cosine similarity with each other. One obvious weakness of the KNN approach is that it is too costly. The time complexity is $O(|E_P||C_P|)$, where $|E_P|$ and $|C_P|$ are the numbers of entities and concepts in Probase respectively, because any entity in Probase is a candidate (overall $O(|E_P|)$ entities) and the cosine similarity computation consumes $O(|C_P|)$ time. Note that Probase contains millions of entities and concepts. Hence, the complexity is generally unacceptable.

In the next section, we propose probabilistic models with higher accuracy and efficiency.

IV. Problem Model

In this section, we propose two models to solve our problem. The first model seeks to maximize a probabilistic based relevance ranking measure. The second aims to minimize the difference between the concept distributions of query entities before and after its acceptance of the suggested entities.

Table II lists the notations we use.

A. A Probabilistic Relevance Model

When given a set of query entities $q = \{q_i | q_i \in E\}$, we model the relevance of an entity e to q with $rel(q, e)$, which can be interpreted as the likelihood that a real person will think of the entity e when he/she observes the entities in the query q. Thus, our objective is to find the entity whose relevance is the highest:

$$\arg\max_{e \in E - q} rel(q, e) \tag{2}$$

Then, the key is to define the relevance function.

Consider the psychological procedure of a user to infer an entity by formulating a set of example entities. A real user tends to formulate the query by referring to some concepts of the examples as well as the target entity. The concepts

referred to describe the one or more aspects of these entities. For example, given {*China, India, Brazil*}, two concepts {*developing country, emerging market*} naturally come to our mind. We define $P(c_i|q)$ to be the typicality that c_i is referred to when we are presented with q. The most typical entity under the concepts the query referred to tends to be recommended. For the example above, in both of the two concepts, *Russia* is a typical entity. Let $P(e|c_i)$ be the typicality of e given concept c_i, which can be computed by Eq. 1. Thus, we have the following ranking function:

$$rel(q, e) = \sum_i P(e|c_i)P(c_i|q) \tag{3}$$

Clearly, the ranking function is positively correlated to the two factors $P(e|c_i)$ and $P(c_i|q)$, which reflect the following two principles: (1) A typical instance should be recommended; (2) If a concept is more likely to be associated with the query entities, its typical instance deserves to be recommended. The summation over all concepts reflects the fact that if there are many concepts leading to an instance e, it should be recommended.

Interpretation of $P(c_i|q)$: The direct interpretation of $P(c_i|q)$ is as follows. Suppose there exists an *ideal* concept for the query entities, which under some circumstances holds true. For example, {*China, India, Brazil*} usually implies that the best concept is *BRIC* countries, which directly helps us find the appropriate entity *Russia*. Thus, $P(c_i|q)$ can also be interpreted as the typicality that c_i is the *ideal* concept of q.

Otherwise, the ideal concept does not always exist in the conceptual taxonomy. For example, it is hard to give an explicit-yet-simple concept to describe {*Baidu, Tecent, Alibaba*}. They refer to the three biggest IT companies in China (BAT for short). In these cases, $P(C|q)$ can be interpreted as the concept distribution of the query entity set. For example, for the query entity set {*China, India, Brazil*}, {*developing country, emerging market*} are two representative concepts to describe the semantics of the query entity set. Thus, each concept c_i can be aggregated with a weight $P(C = c_i|q)$ (for short $P(c_i|q)$), to represent the semantics of q together.

However, some concepts are too general to be the ideal concepts to suggest other entities. In our running example, *country* is a concept which can summarize all the query examples but it makes little difference in evaluating other entities of countries. These concepts are relatively vague to characterize the given entities. Thus, we introduce an item $\delta(c_i)$ to help choose the ideal concepts given q. Then our new ranking function will be:

$$rel(q, e) = \sum_i P(e|c_i)P(c_i|q)\delta(c_i) \tag{4}$$

We propose two strategies to compute $\delta(c_i)$, one is to penalize the popular concepts based on the entities of the concepts, the other one is to find fine-grained concepts based on the hierarchical structure of the conceptual taxonomy. We will elaborate them as well as the computation of $P(c_i|q)$ in Section V.

B. A Relative Entropy Model

An alternative model is to use a concept distribution of query entities. We find the e such that its admission into q has the least impact on the original concept distribution. More formally, $P(C|q)$ is the concept distribution given query entity set q, which can be represented as a set of vectors $\{< c_i, P(c_i|q) >\}$. We just need to find the entity e such that $P(C|q, e)$ is closest to $P(C|q)$. A popular measure of the distance between two probability distributions is KL-divergence, also known as relative entropy. Thus, our problem can be modelled as:

$$\arg\min_{e \in E-q} KL(P(C|q), P(C|q, e)) \tag{5}$$

where KL-divergence is defined as:

$$KL(P(C|q), P(C|q, e)) = \sum_{i=1}^{n} P(c_i|q) \times log(\frac{P(c_i|q)}{P(c_i|q, e)}) \tag{6}$$

KL-divergence between $P(C|q)$ and $P(C|q, e)$ characterizes the expectation of the logarithmic difference between the two probability distributions. Note that KL-divergence is not symmetric, which means that $KL(P(C|q), P(C|q, e)) \neq KL(P(C|q, e), P(C|q))$. The rationality of our optimization function is that the posterior typicality of c_i observing q (i.e. $P(c_i|q)$) is more confident than observing q and an arbitrary entity e. Furthermore, as the same reason described in section IV-A, Eq. 6 should also carefully choose the concepts by evaluating $\delta(c_i)$ and $P(c_i|q)$, which will be elaborated in the next section.

Next, we justify why minimizing the distribution disparity in our problem setting is an effective objective.

Rationality of the objective function: We justify the minimization objective by a hypothesis test. The basic idea is to show that the admission of a right entity will preserve the concept distribution, with statistical significance, under the comparison to a null model. Specifically, for each query entity set q with the ground truth, we choose the answer entity e to compute $d_1 = KL(P(C|q), P(C|q, e))$, and randomly choose another entity $e' \neq e$ which shares at least one concept with any query entity in q to compute $d_2 = KL(P(C|q), P(C|q, e'))$. We repeat the random selection 50 times and summarize the average of d_2, denoted by $\bar{d_2}$. Thus, for query entity set q, we have a pair of distances $(d_1, \bar{d_2})$. We calculate these paired distances for 50 different q.

Then, we form the null hypothesis as: *there is no difference between d_1 and $\bar{d_2}$*. We test the hypothesis on all the pairs of distances by *paired t-test*. The result shows that the P-value score is less than 0.001, which is much smaller than the significance level 0.05. The results suggest that there is sufficient statistical evidence to reject the null hypothesis. In other words, the admission of the answer entity leads to a smaller distribution distance than a randomly selected one.

Figure 1 shows that our observation above consistently holds for six cases listed in Table I.

V. Model Computation

In this section, we elaborate how to compute $\delta(c_i)$ and $P(c_i|q)$ which are two common parts of the two probabilistic

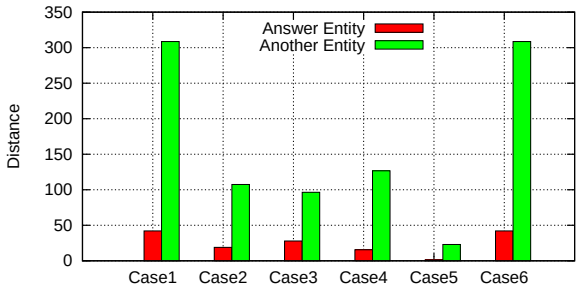

Fig. 1: Case study on the objective function

models that we have not discussed in detail. We propose two approaches for $\delta(c_i)$ and $P(c_i|q)$ respectively .

A. $\delta(c_i)$ Computation

As we have mentioned in Section IV, $\delta(c_i)$ evaluates the concepts which will be used as a latent variable to suggest entities. The basic idea is to find fine-grained underlying concepts for queries instead of those which are too general or too specific. A concept too general may be very likely to introduce noisy entities leading to low precision, while a concept too specific may be difficult to find related entities leading to low recall. Therefore, we propose two approaches here, the first one is to directly penalize popular concepts based on the entities of the concept, the second one is to use a granularity-aware model to find fine-grained concepts based on the hierarchical structure of the conceptual taxonomy .

1) Penalizing Popular Concepts: From the perspective of entities to observe concepts, a concept with more entities may be more general. For example, in Probase, *country* has 2648 entities, *developing country* has 149 entities, while *growing market* only has 18 entities. It is obvious that *country* is a concept more general than *developing country* and *growing market*, which is rational according to the institution of human beings. Thus, some concepts may have high $P(c_i|q)$ due to their popularity, such as *company* for $\{$*Tencent, Baidu*$\}$. These concepts are relatively vague to characterize the given entities because they also contain many other entities and may even lead to the issue of semantic drift. It motivates us to introduce the typicality of the concept $P(c_i)$ to penalize a popular concept.

Then, $\delta(c_i)$ will be:

$$\delta(c_i) = \frac{1}{P(c_i)} \tag{7}$$

where $P(c_i)$ is defined as follows:

$$P(c_i) = \frac{n(c_i)}{\sum_{c \in C} n(c)} \tag{8}$$

where $n(c_i)$ is the number of occurrence of c_i derived from Probase data.

2) Granularity-aware Approach: When using $P(c_i)$ to penalize popular concepts, a potential problem may be bias to the concepts which are too specific which may lead to more difficulty in finding related entities since specific concepts always contain few entities. Thus, we propose a new method considering granularity of the concepts by using the measure of hitting time to the query based on the hierarchy of the concepts. We first give the computation method and then state the rationality.

Let C_q^k be a set with k fine-grained concepts underlying query q, the expected hitting time $H(q|C_q^k)$ of the random walk is the sum of the expected number of steps before each $q_i \in q$ is visiting a concept c in C_q^k.

Thus, our target becomes finding a set C_q^k such that:

$$\arg\min_{C_q^k} H(q|C_q^k) \tag{9}$$

Here k can control the number of the concepts we choose, and the larger the k is, the more general concepts will be introduced into the set, and to be used to suggest entities.

Naturally, $H(q|C_q^k)$ can be computed as follows:

$$H(q|C_q^k) = \sum_{q_i \in q} \sum_{c \in C_q^k} h(q_i|c) \tag{10}$$

where $h(q_i|c)$ means the expected number of steps before a query entity q_i is visiting a concept c.

It can be easily verified that the hitting time satisfies the following system of linear equations [29]:

$$\begin{cases} h(q_i|c) = 0, & if q_i = c \\ h(q_i|c) = 1 + \sum_{c' \in c(q_i)} P(c'|q_i) h(c'|c), & if q_i \neq c \end{cases} \tag{11}$$

Here we use $P(c'|q_i)$, the typicality of observing concept c' when given q_i to be the probability of starting from q_i to c.

Then, $\delta(c_i)$ will be:

$$\begin{cases} \delta(c_i) = 1, & if c_i \in C_q^k \\ \delta(c_i) = 0, & if c_i \notin C_q^k \end{cases} \tag{12}$$

The rationality of the model: As for our running example, we can find that both *country* and *developing country* related to all entities in the query. However, it is obvious that *country* is a concept which is too general to be the underlying concept when given $\{$*China, India, Brazil*$\}$, and *developing country* may be a better one. It is easy to be observed that *country* is also a concept of *developing country*. That means there may exist some longer paths which will increase the expected steps starting from the query entity to the general concept. Thus, a concept not too general should be a concept which has a short expected distance to query entities. To avoid a concept too specific, the concept should have a short expected distance to every entities in the query set q. Thus, the optimization target should be to find k concepts which can minimize $H(q|C_q^k)$.

Note that to make our solution efficient, two strategies can be used. First, $h(e|c)$ can be computed offline which will not increase the computation time when given queries to suggest entities. Second, since we only interested in concepts within

a short hitting time, we may just ignore the concepts with a hitting time larger than a certain threshold of steps which can dramatically shorten the computation cost.

B. $P(c_i|q)$ Computation

In this part, we elaborate how to compute $P(c_i|q)$ which evaluate the extent of the concept underlying the query entities.

1) Naive Bayes Model: We first propose a Naive Bayes approach to compute $P(c_i|q)$. According to the Bayes theorem, we have:

$$P(c_i|q) = \frac{P(q|c_i)P(c_i)}{P(q)} \propto P(q|c_i)P(c_i) \tag{13}$$

Since $P(q)$ is only dependent on the query, it can be ignored for the purpose of ranking.

In general, a person's choices of two entities e_i, e_j are logically independent with each other when given concept c_i. Thus, we can have an independence assumption that:

$$P(e_j, e_k|c_i)_{\forall e_j, e_k \in q} = P(e_j|c_i)P(e_k|c_i) \tag{14}$$

Then, we have:

$$P(c_i|q) \propto \prod_{e_j \in q} P(e_j|c_i)P(c_i) \tag{15}$$

Generally, there are relatively few concepts related to all of the query entities, therefore appropriate smoothing is necessary to avoid zero probabilities. To do this, we can assume that with probability λ, the user would choose the concept by its prior typicality. Thus, we can rewrite the Eq.15:
$P(c_i|q) \propto$

$$P(c_i) \prod_{e_j \in q, n(e_j,c_i)>0} \lambda P(e_j|c_i) \prod_{e_j \in q, n(e_j,c_i)=0} (1-\lambda)P(e_j) \tag{16}$$

where $P(e_j|c_i)$ can be computed by Eq. 1, and $n(e_j, c_i)$ is the number of co-occurrences between e_j and c_i in Probase. The prior typicality of $P(c_i)$ and $P(e_j)$ can be computed by the following equation:

$$P(c_i) \propto n(c_i) \tag{17}$$

$$P(e_j) \propto n(e_j) \tag{18}$$

where $n(c_i)$ and $n(e_j)$ is the number of occurrence of c_i and e_j in Probase.

Note that when using this model in the relative entropy model, $P(c_i|q, e)$ depends on $P(q, e)$, which can not be ignored. Thus, the computation of $\frac{P(q,e)}{P(q)}$ in Eq. 6 can be simplified to be $P(e)$ by the independence assumption.

2) A Noisy-Or Model: Alternatively, we can mimic a psychological process of identifying an ideal concept, when query instances are presented one by one to a human—As more query entities are given, desirable concepts will amplify and eventually peak, as illustrated in Figure 2. This observation implies that the signal indicating the right concept should be amplified when more entities are observed, and the single indicating the incorrect concept should be weakened. All these observations motivate us to use a Noisy-Or model to compute $P(c_i|q)$:

$$P(c_i|q) = 1 - \prod_{e_j \in q} (1 - P(c_i|e_j)) \tag{19}$$

where $P(c_i|e_j) = \frac{n(c_i,e_j)}{n(e_j)}$ can be computed by Probase data.

Given $P(c_i|q)$, we can use it to compute $P(c_i|q, e)$ incrementally, which avoids the wasteful computation.

$$\begin{aligned} P(c_i|q, e) &= 1 - \prod_{e_j \in q} (1 - P(c_i|e_j))(1 - P(c_i|e)) \\ &= 1 - (1 - P(c_i|q))(1 - P(c_i|e)) \\ &= P(c_i|q) + P(c_i|e) - P(c_i|q)P(c_i|e) \end{aligned} \tag{20}$$

where $P(c_i|q)$ is given and the other two parts can be computed by Probase.

Rationality of Noisy-Or Model: The Noisy-Or model reflects the following rationality: *the concept covering more query entities has a higher conditional probability*, which is consistent with our intuitions. We illustrate this by the following example. Still given {*China, India, Brazil, Russia*}, when we compute the query entities' distribution of their 20 concepts one by one, we can observe the change in Figure 2. It realized that the conditional probability became higher and concentrated to certain concepts. We also illustrate this in Example 1.

Example 1 (Rationality of Noisy-Or Model): Given query entity set {*China, India, Brazil*}, if {*China, India, Brazil*} are all related to the concept {*developing country, market*}, it's obvious that these two concepts are both very important. The concept *Latin American country* is only related to the entity *Brazil*, so it will be less important to the concept {*developing country, market*}. Thus, the more query entities the concept is related to, the more important the concept is. Furthermore, the signal indicating the incorrect concept will be weakened.

VI. Experiment Evaluation

In this section, we systematically evaluate the effectiveness of our models and solutions with the state-of-the-art approaches on two ground truth data sets.

The first one is the data set presented in Wang and Cohen [16] which are simple conceptual lists. To highlight the advantage of our approaches, using fine-grained underlying concepts, we constructed a tougher data set to test our solutions. The second data set consists of lists that are deliberately selected from Wikipedia. These lists can only be described by complicated (fine-grained) concepts and might not have explicit simple concept names, and we call them complicated-yet-typical concepts. This data set is much more challenging because it is more likely to introduce false-positives because of the granularity of the concepts underlying the example entities. Then, we introduce our two data sets in detail and present the experimental results on each data set respectively.

We evaluate the effectiveness of our two models (PRM and REM) with two computation methods of $\delta(c_i)$ and $P(c_i|q)$ respectively, Penalizing popular concepts (PP for short), Granularity-aware approach (FG for short), Naive

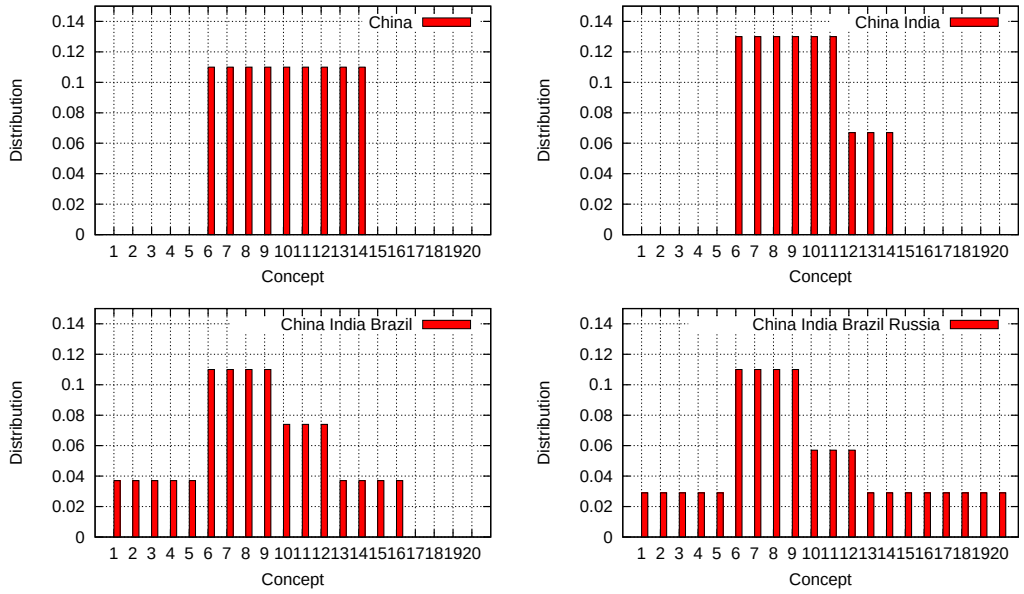

Fig. 2: Case study on the change of the concept distribution

Bayes approach (BA for short) and Noisy-Or model (NO for short). Thus, we have overall 8 versions of our solution: PRM+PP+BA, PRM+FG+BA, PRM+PP+NO, PRM+FG+NO, REM+PP+BA, REM+FG+BA, REM+PP+NO, REM+FG+NO. We compared them with the baseline KNN method we proposed, the entire list completion proposed in [30] (ER for short), the structure-based approach using properties of entities in [7] (ESBA for short) and SEISA [14], one of the strongest baselines [25]. Unluckily, since we do not have the web list data used in SEISA, we re-implemented it by replacing the web lists with the concepts and their entities in Probase.

A. Simple Conceptual Lists

Set up: We use lists in English in the seal data set [16] as our first test data set. We randomly choose two instances of each list as the query and evaluate the quality of the ranked result lists returned by different solutions.

Competitors: We compared our results with KNN, ESBA, SEISA and the results of ER reported in [30]. Since ER has different versions, we report the performance results derived by the version that P (or R, F) is best. We denote these versions by ER_p, ER_r, ER_f, respectively.

Note that here we implement our approaches with granularity-aware approach with $k = 200$ because in the first data set, concepts are all general concepts, thus we need to set a relatively large k, here 200, and we will study the effect when varying k from different numbers.

Metrics: We use precision, recall and F-score as the metrics on our data set. Let the number of the entities in the ranked list be n_R and the number of the entities in the seal list be n_L. Thus, precision P is the number of correct entities that are in the ranked list divided by n_R. Recall R is the number of the correct entities which appear in the ranked list, divided by n_L. F-scores are the harmonic mean of the recall and precision.

Results:

Precision, Recall and F-score: The comparison results in Table III show that almost all of our solutions consistently outperform the competitors on all the tested measures except the precision of REM+BA. It sufficiently suggests that the conceptual taxonomy is beneficial for the entity retrieval task in most cases.

The detailed comparisons reveal that PRM+FG+NO has the highest precision, and REM+FG+NO has the highest recall. When comparing our approaches with FG and PP, BA and NO respectively, we can find that the methods with FG are better than the ones with PP in most cases, and the methods with NO are always better than the ones with BA. It reveals that the Noisy-Or model computes the conditional concept distribution more accurately than the Naive Bayes Model, and very often the methods that directly penalize popular concepts may make some specific concepts greatly affect the results and make it more difficult to find related entities.

Influence of # concepts underlying the query: In Figure 3 and 4, we study the effect of the number of the concepts underlying the query examples on the results. As we have mentioned, the larger the number is, the more general concepts will be introduced into the concept distribution. However, here we can see that the number do not make a big difference. The reason is very likely to be the fact that the query examples generated by the lists of the seal data set are all from general concepts, thus the fine-grained concepts may not play an important role in generating the results.

B. Fine-grained Conceptual Lists

As we have noticed, the lists in seal data set have simple concept names, such as *us-presidents*. *Us presidents* is a specific and relatively simple concept, whose counterpart can be easily found in Probase. Thus, our task is reduced to the retrieval of entities from the concepts in Probase. This fact implies that our

TABLE III: Results on the first dataset

Method	P	R	F
KNN	0.13	0.230	0.166
ESBA	0.221	0.263	0.241
SEISA	0.240	0.286	0.259
PRM+PP+BA	0.257	0.363	0.317
PRM+FG+BA	0.294	0.449	0.355
PRM+PP+NO	0.310	0.409	0.359
PRM+FG+NO	**0.380**	0.537	**0.445**
REM+PP+BA	0.180	0.428	0.253
REM+FG+BA	0.196	0.422	0.268
REM+PP+NO	0.280	0.444	0.364
REM+FG+NO	0.362	**0.570**	0.443
ER_p	0.227	0.142	0.137
ER_r	0.121	0.236	0.136
ER_f	0.204	0.209	0.158

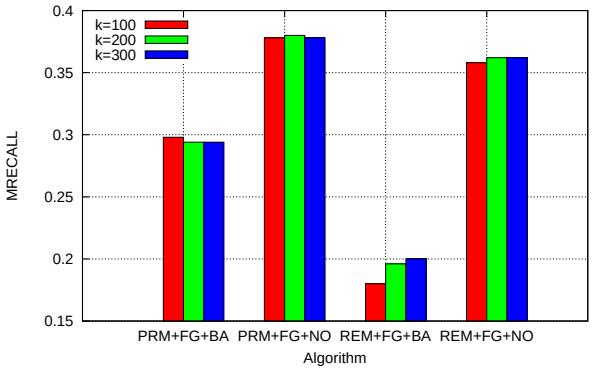

Fig. 3: MAP of granularity-aware approach varying k

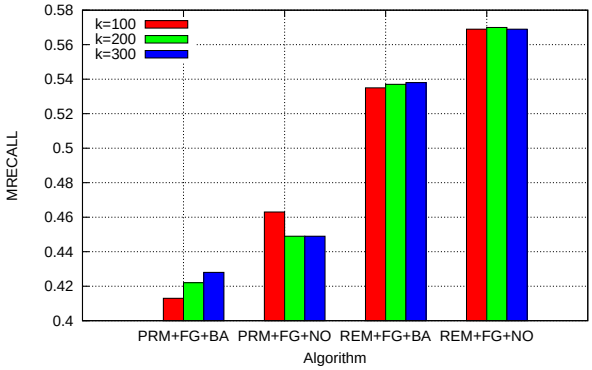

Fig. 4: MRECALL of granularity-aware approach varing k

idea to use conceptual taxonomy is rational. On the other hand, this data set somewhat biases towards our solutions. Hence, it motivates us to construct a tougher data set, which sufficiently reveals the advantage of our models and solutions, specifically our approaches can find the entities of fine-grained conceptual lists with high precision.

For this purpose, we deliberately select some complicated-yet-typical concepts from the Wikipedia articles. Most of these concepts have the name such as *Big N* and *Great N*. These article pages contain a list of entities that share the same specific concept, which usually requires a complicated description. For example, from the *Big 4 (tennis)* page in Wikipedia, we can get four entities {*Roger Federer, Rafael Nadal, Novak Djokovic, Andy Murray*}. The full description of the concept actually is *the four most famous tennis players in the world nowadays*. We collected 112 such lists to construct the second data set. Some example lists as well as their complicated concept descriptions are shown in Table IV.

Competitors: We compared our solutions on the second data set with KNN, ESBA, SEISA and ER proposed in [30] which implemented the entity ranking task with Bayesian Inference using Wikipedia data. To compare with it fairly, we re-implemented its method on Probase data. We used the concept in Probase as the word item (the hidden random variable θ) in Wikipedia, and directly computed $P(e|\theta)$ (the conditional probability of entity e given θ), $P(\theta|D)$ (the conditional probability of θ given the example entities D), and $P(e)$ based on the probability in Probase. We denote this approach as ER+BS.

Query construction: Given the ground truth data, we use the following way to construct the query set for the evaluations:

For each list L with $|L|$ instances, we use $\sigma(L) = \lceil \alpha |L| \rceil$ entities as the query entities. Ideally, we hope our solution can find the remaining $1 - \alpha$ entities. In our experiment, we choose α as $\frac{1}{2}$, $\frac{2}{3}$, $\frac{3}{4}$, and we will also study its influence on the results.

Metrics: For each query, the answer entities should be ranked higher than other unrelated entities. Thus, on this data set, we use average NDCG to evaluate each query in Q, denoted as $mndcg$. Obviously, a larger $mndcg$ implies a better ranking.

Results:

mndcg: The experimental results on the complicated conceptual lists are shown in Figure 5. It is evident that except REM+BA, our approaches are all better than the baselines. We can find that REM+FG+NO has the highest mndcg even varying the number of the examples. Comparing the results in detail, it is obvious that FG+NO are always better than PP+BA, which reveals two conclusions that first, Noisy-Or model is more accurate in modeling the conditional probability when given some example entities. Second, granularity-aware approach is better in suggesting entities underlying fine-grained concepts.

In Figure 6, we study the precision@k of different methods. We can find that our method REM+FG+NO has a relatively higher precision than other methods even when k is small, and its precision can arrive at above 80% when k is 3. Furthermore,

TABLE IV: Fine-grained conceptual lists

Entity List	The full description of the concept
Leonnado da vinci, Raphael, Michelangelo	the three most famous art masters of the renaissance
Aaron Kwok, Jacky Cheung, Leon Lai, Andy Lau	the four most famous singers in hongkong
Cats, Miss Saigon, Les Miserables, Phantom of the Opera	the four most famous and classical musicals
Agricultural Bank of China, China Construction Bank, Bank of China, Industrial and Commercial Bank Of China	the four biggest banks owned by chinese government

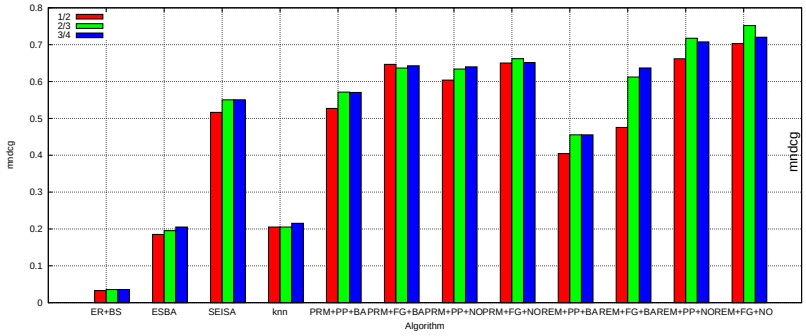

Fig. 5: Performance of different solutions

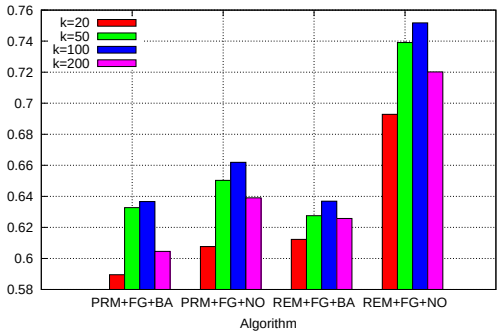

Fig. 7: Varying k of granularity-aware method

PRM+NO models also perform well when k is small which reveals the rationality of our models.

Influence of # concepts underlying the query : We study the choice of the number of the concepts underlying the query examples. We vary k from 20, 50, 100 to 200 to observe the results. We can find that when we choose 50 or 100, the $mndcg$ may be higher than we choose 20 or 200. It is reasonable because 20 fine-grained concepts may be a little too specific to find accurate related entities, while 200 concepts may introduce more general concepts leading to incorporating many false positive entities. Therefore, k, the number reflecting the granularity should be carefully chosen.

C. Case studies

In this subsection, we give case studies to show the rationality of our models and solutions.

Table V and VI presents studies on some cases from our second data set. The query examples are in the second row of the tables, and the entities in the third row are answer entities which complete the complicated-yet-typical concepts where the example entities come from. They are possible entities which people may want the computer to return when given the query entities. Of course, some other entities are also very likely to be associated when given these entities, but here we consider that the result lists which contain these possible answers have a higher quality.

We present the top-5 entities of the result lists. Because of the space limitation, we show the results of the best competitor and the best approach of ours. In all of the cases, the possible answers appear in the top of the lists of our results. While the baseline tend to introduce false positive entities, and give them higher ranks. Concretely, in Case 1, *alibaba* and *tencent* are Chinese internet giants, while SEISA incorporated a lot of companies in the U.S. It is very likely that SEISA used the concept *company* leading to the false positive entities. The same situation also happened in Case 2, and Case 3. In Case 4, SEISA seemed to include some unrelated entities.

To testify the contribution of the fine-grained concepts, in Figure VII, we show the latent results of entity suggestion. Here, we show top-5 concepts mined by our method. We can see that most of our concepts are related and not too general or specific.

The above results sufficiently show that the biased selection of concepts is critical for the effectiveness of entity inference.

VII. Conclusion

This paper studies entity suggestion by example, a technique with diverse applications. Though many solutions exist, they are built mostly on the idea of co-occurrence in the text or web lists and exhibit limited effectiveness. In contrast, we leverage the new opportunities brought by many web-scale conceptual taxonomies with rich isA concept-instance relationships. Specifically, we have proposed two probabilistic approaches, the first leveraging the typicality of concepts and entities to make the inference biased toward the more promising entities, and the second using an optimization solution to minimize the difference before and after the acceptance of a candidate entity. With these two models, we have solved the challenging problems of how to aggregate the conceptual information of the example entities by a Naive Bayes Model

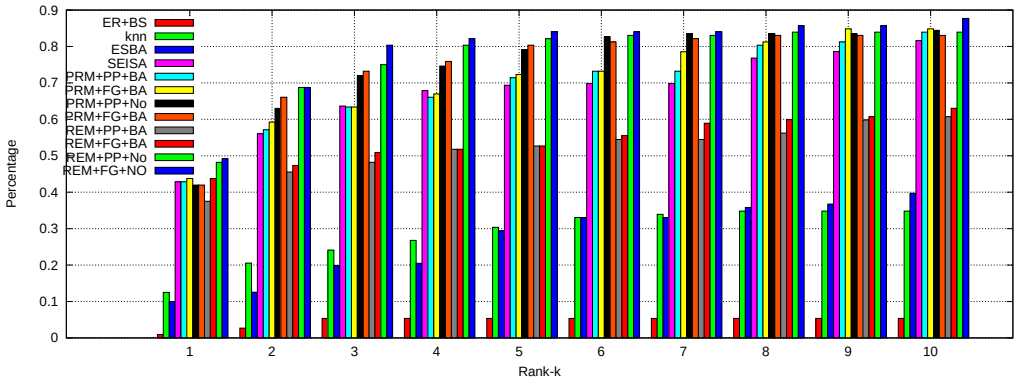

Fig. 6: Precision@K of different methods

TABLE V: Top-5 results of four cases

Case 1		**Case 2**	
Query: alibaba, tencent		**Query:** agriculture bank of china, china construction bank	
Possible Answer: baidu		**Possible Answer:** bank of china, industrial and commercial bank of china	
REM+FG+NO	SEISA	REM+FG+NO	SEISA
baidu	ebay	icbc	bank of china
neusoft group	baidu	bank of china	bank
alibaba.com	amazon	china development bank	bank of america
lenovo group	facebook	bank of communications	commercial bank of china
byd	global source	china international capital corporation limited	industrial

TABLE VI: Top-5 results of four cases

Case 3		**Case 4**	
Query: aaron kwok, jacky cheung		**Query:** kmpg, pricewaterhousecoopers	
Possible Answer: leon lai, andy lau		**Possible Answer:** ernst, deloitee	
REM+FG+NO	SEISA	REM+FG+NO	SEISA
andy lau	madonna	ernst & young	ibm
leon lai	kanye west	deloitte	accenture
edison chen	bette midler	ernst	fidelity investments
george guilder	beyonce	microsoft	schwab
george shea	aretha franklin	pwc	acs

and a Noisy-Or model, how to find fine-grained concepts by an entity-based approach and a hierarchy-based approach and how to avoid the false positive in the inference due to the absence of contextual information. We have validated the effectiveness of our approaches using extensive evaluations with real-life data.

REFERENCES

[1] B. Dalvi, J. Callan, and W. Cohen, "Entity list completion using set expansion techniques," DTIC Document, Tech. Rep., 2011.

[2] D. E. Rose and D. Levinson, "Understanding user goals in web search," in *Proceedings of the 13th international conference on World Wide Web*. ACM, 2004, pp. 13–19.

[3] R. Delbru, S. Campinas, and G. Tummarello, "Searching web data: An entity retrieval and high-performance indexing model," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 10, pp. 33–58, 2012.

[4] E. Meij, K. Balog, and D. Odijk, "Entity linking and retrieval for semantic search," in *Proceedings of the 7th ACM international conference on Web search and data mining*. ACM, 2014, pp. 683–684.

[5] X. Yu, "Entity recommendation and search in heterogeneous information networks," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2014.

[6] K. Balog, M. Bron, and M. De Rijke, "Query modeling for entity search based on terms, categories, and examples," *ACM Transactions on Information Systems (TOIS)*, vol. 29, no. 4, p. 22, 2011.

[7] M. Bron, K. Balog, and M. De Rijke, "Example based entity search in the web of data," in *Advances in Information Retrieval*. Springer, 2013, pp. 392–403.

[8] K. Balog, P. Serdyukov, and A. P. d. Vries, "Overview of the trec 2010 entity track," DTIC Document, Tech. Rep., 2010.

[9] D. Mottin, T. Palpanas, and Y. Velegrakis, "Entity ranking using click-log information," *Intelligent Data Analysis*, vol. 17, no. 5, pp. 837–856, 2013.

[10] M. M. Zloof, "Query by example," in *Proceedings of the May 19-22, 1975, national computer conference and exposition*. ACM, 1975, pp. 431–438.

[11] K. Ahn, J. Bos, D. Kor, M. Nissim, B. L. Webber, and J. R. Curran, "Question answering with qed at trec 2005." in *TREC*, 2005.

[12] D. Mottin, M. Lissandrini, Y. Velegrakis, and T. Palpanas, "Exemplar

TABLE VII: Top-5 concepts

Case 1	Case 2	Case 3	Case 4
chinese internet giant	chinese bank	popular canto-pop entertainer	accounting firm
shanghai-chinese internet giant	mainland bank	hot hongkong singer	global consultant
B2b service provider	chinese government bank	hongkong artists	international accounting firm
chinese consumption stock	chinese leader	universal's other frontline artist	big accounting firm
chinese social networking site	bank	spokesperson	global business consultancy firm

queries: Give me an example of what you need," *Proceedings of the VLDB Endowment*, vol. 7, no. 5, 2014.

[13] X. Gao, B. Xiao, D. Tao, and X. Li, "A survey of graph edit distance," *Pattern Analysis and applications*, vol. 13, no. 1, pp. 113–129, 2010.

[14] Y. He and D. Xin, "Seisa: set expansion by iterative similarity aggregation," in *Proceedings of the 20th international conference on World wide web*. ACM, 2011, pp. 427–436.

[15] G. L. Murphy and D. L. Medin, "The role of theories in conceptual coherence." *Psychological review*, vol. 92, no. 3, p. 289, 1985.

[16] R. C. Wang and W. W. Cohen, "Language-independent set expansion of named entities using the web," in *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*. IEEE, 2007, pp. 342–350.

[17] R. Blanco, B. B. Cambazoglu, P. Mika, and N. Torzec, "Entity recommendations in web search," in *The Semantic Web–ISWC 2013*. Springer, 2013, pp. 33–48.

[18] S. Metzger, R. Schenkel, and M. Sydow, "Aspect-based similar entity search in semantic knowledge graphs with diversity-awareness and relaxation," in *Proceedings of the 2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)-Volume 01*. IEEE Computer Society, 2014, pp. 60–69.

[19] N. Jayaram, M. Gupta, A. Khan, C. Li, X. Yan, and R. Elmasri, "Gqbe: Querying knowledge graphs by example entity tuples," in *Data Engineering (ICDE), 2014 IEEE 30th International Conference on*, March 2014, pp. 1250–1253.

[20] Google, "Googlesets." URL:http://labs.google.com/sets, accessed on 04-10-2006, 2006.

[21] Z. Ghahramani and K. Heller, "Bayesian sets," 2005.

[22] R. C. Wang and W. W. Cohen, "Iterative set expansion of named entities using the web," in *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*. IEEE, 2008, pp. 1091–1096.

[23] L. Sarmento, V. Jijkuon, M. de Rijke, and E. Oliveira, "More like these: growing entity classes from seeds," in *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*. ACM, 2007, pp. 959–962.

[24] P. Pantel, E. Crestan, A. Borkovsky, A.-M. Popescu, and V. Vyas, "Web-scale distributional similarity and entity set expansion," in *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2-Volume 2*. Association for Computational Linguistics, 2009, pp. 938–947.

[25] C. Wang, K. Chakrabarti, Y. He, K. Ganjam, Z. Chen, and P. A. Bernstein, "Expansion of tail concept using web tables," 2014.

[26] D. Kim, H. Wang, and A. Oh, "Context-dependent conceptualization," in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 2013, pp. 2654–2661.

[27] Y. Song, H. Wang, Z. Wang, H. Li, and W. Chen, "Short text conceptualization using a probabilistic knowledgebase," in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Three*, ser. IJCAI'11. AAAI Press, 2011, pp. 2330–2336. [Online]. Available: http://dx.doi.org/10.5591/978-1-57735-516-8/IJCAI11-388

[28] W. Wu, H. Li, H. Wang, and K. Zhu, "Towards a probabilistic taxonomy of many concepts," Tech. Rep., 2011.

[29] F. Fouss, A. Pirotte, J.-M. Renders, and M. Saerens, "Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation," *Knowledge and data engineering, ieee transactions on*, vol. 19, no. 3, pp. 355–369, 2007.

[30] S. F. Adafre, M. de Rijke, and E. T. K. Sang, "Entity retrieval," *Recent Advances in Natural Language Processing (RANLP 2007)*, 2007.